Reconfigurable Antennas, Preemptive Switching and Virtual Channel Management

Santhosh Kumar, Jean-Francois Chamberland, Gregory H. Huff
Department of Electrical and Computer Engineering,
Texas A&M University

Abstract—This article considers the performance of wireless communication systems that utilize reconfigurable or pattern-dynamic antennas. The focus is on finite-state channels with memory and performance is assessed in terms of real-time behavior. In a wireless setting, when a slow fading channel enters a deep fade, the corresponding communication system faces the threat of successive decoding failures at the destination. Under such circumstances, rapidly getting out of deep fades becomes a priority. Recent advances in fast reconfigurable antennas provide new means to alter the statistical profile of fading channels and thereby reduce the probability of prolonged fades. Fast reconfigurable antennas are therefore poised to improve overall performance, especially for delay-sensitive traffic in slow-fading environments. This potential for enhanced performance motivates this study of the temporal behavior of point-to-point communication systems with reconfigurable antennas. Specifically, agile wireless communication schemes over erasure channels are analyzed; situations where using reconfigurable antennas yield substantial performance gains in terms of throughput and average delay are identified. Scenarios where only partial state information is available at the receiver are also examined, naturally leading to partially observable decision processes.

Index Terms—Data communication, Wireless communication, Antennas, Reconfigurable architectures, Adaptive systems Communication channels.

I. INTRODUCTION

High-performance reconfigurable antenna technologies are collectively emerging as a viable option for smartphones, mobile hotspots, and portable computers [1], [2], [3], [4], [5]. Reconfigurable antennas can be employed in a number of ways to increase the connectivity profile and robustness of wireless communication channels, and they create new and promising opportunities for the engineering of superior communication schemes. Such antennas are designed to intentionally and reversibly alter the character of their performance-governing electromagnetic fields. As a result, they are able to modify the directional and polarization properties of their radiation patterns and thereby change the spatiotemporal characteristics of the communication channels they induce.

Beyond the basic properties of reconfigurable antennas, it is interesting to note that manipulations of radiation patterns can be automated through rapid feedback and triggered by events such as deep fades and successive decoding failures

[6], [7], [8], [9]. This enables the implementation of closed-loop systems with the cognitive ability to seamlessly adapt to evolving electromagnetic environments and interference conditions. The capacity to provide these features in (near) real time is primarily determined by the speed and complexity of the reconfiguration mechanisms used to facilitate electromagnetic agility within the reconfigurable antenna structures. This highlights a natural tradeoff between the potential benefits of a configuration change and the downtime associated with each transformation event.

Although reconfigurable antennas have been and continue to be the subject of concerted research efforts in the antennas and propagation community, a detailed analysis of their repercussions on the foundations of wireless communications is still lacking. Key to the widespread adoption of such technologies, aside from miniaturization, is provable gains in terms of capacity, delay-throughput profile and network connectivity [10], [11]. This article seeks to address the need for better understanding the impact of adaptive antenna systems. From a conceptual point of view, fast reconfigurable antennas can be employed to establish ancillary virtual links between two devices. As the quality of the current channel degrades, it may become advantageous to transition to an alternate antenna state and, consequently, to another channel realization.

This added flexibility at the physical layer is likely to boost the perceived performance of delay-sensitive applications over channels with memory. Indeed, several studies document the fact that channel variations are particularly detrimental to delay-constrained communications [12], [13], [14], [15]. Channel memory further exacerbates this situation, as it increases the propensity for prolonged deep fades [16], [17], [18]. Thus, having the capability to jump to a different virtual channel seems an attractive option in these circumstances.

To carry our analysis of reconfigurable systems, we leverage several results that have appeared in the literature in the past. First, we adopt a class for channels with memory similar to the finite-state channel model proposed by Gilbert and Elliott [19], [20]. The delay-sensitive aspect of the problem is captured through a queueing formulation whose solution is obtained, partly, by applying techniques originally developed by Neuts [21], [22], [23]. The conceptual bridge between the physical layer and the queueing system is provided by error correcting codes and the availability of feedback. The framework presented below parallels some of our previous work [17], [18]. The incorporation of reconfigurable antennas into the problem setting and the insights provided by our

analysis are novel. This sheds new lights on the potential benefits of adaptive antenna systems and their application to delay-sensitive communications over wireless channels.

Special attention is given to the realistic scenario where the channel state is not known perfectly at the receiver. Rather, it must be inferred from available observations. In this latter case, the optimization problem becomes a partially observable decision process. This forces the controller to make decisions based on an estimated distribution over the channel state. This problem is cast into a dynamic programming framework with limited state knowledge. Known system dynamics and partial observations are combined at every step to update the estimated distribution and, subsequently, take action.

The remainder of this article is organized as follows. The system components, along with a mathematical abstraction for reconfigurable antennas, are described in Section II. Two modes of operation are considered, a classical system with a static antenna structure and an adaptive implementation with the ability to reconfigure the RF front-end depending on the channel state. The evolution of these two competing schemes is characterized in Section III, where we analyze their queueing behaviors. This gives rise to various performance criteria, including throughput, mean waiting time and the probability of the queue exceeding a certain threshold. Pertinent numerical examples are presented in Section V. Finally, concluding remarks and avenues of future research are discussed in the last section.

II. SYSTEM MODEL

In this study, we consider one side of the communication process. That is, information flows from a transmitter to a destination. An eventual assumption on the availability of feedback will necessarily imply the presence of a reverse link. Nonetheless, for the sake of simplicity, we focus on a single direction with the understanding that a similar analysis can be applied to the reverse link. We keep the descriptions of the system model to a minimum and refer the reader to [17], [18] for an elaborate exposition.

As mentioned above, overall system performance is assessed using a queueing formulation. The state of the queue at the source is governed by arrivals and departures. The evolution of the queue is modeled as a discrete-time stochastic process that is synchronized with codeword transmissions. Specifically, during each codeword cycle, we assume that a data packet arrives with probability γ, independently of other time instants. The number of information bits per data packet, denoted by L, is also random and possesses a geometric distribution with parameter ρ. We emphasize that this arrival process is carefully selected to facilitate analysis. Since this article is primarily concerned with identifying the relative benefits of reconfigurable antennas over traditional implementations, adhering to a specific arrival profile is unlikely to affect the nature of our subsequent results.

A. Traditional Implementation

We begin our discussion of the queueing system with a mathematical description of the channel model associated
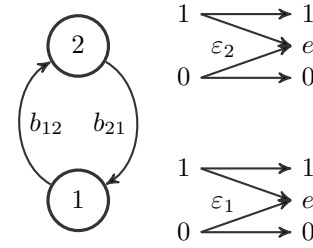

Fig. 1. A finite-state channel with memory is used to model the operation of a wireless link at the bit level. For illustrative purposes, the channel is depicted with only two states, a form known as the Gilbert-Elliott channel.

with a fixed antenna configuration. This communication channel can operate in one of several modes, designated by $\mathcal{C} = \{1, \ldots, k\}$. The evolution of the channel forms a time-homogeneous Markov process and its probability transition matrix is denoted by

$$\mathbf{B} = \begin{bmatrix} b_{11} & b_{12} & \cdots & b_{1k} \\ b_{21} & b_{22} & \cdots & b_{2k} \\ \vdots & \vdots & \ddots & \vdots \\ b_{k1} & b_{k2} & \cdots & b_{kk} \end{bmatrix}, \tag{1}$$

where b_{ij} symbolizes the probability of jumping to state j given that the chain is currently in state i. That is, $\mathbf{B}$ is a right stochastic matrix. When in state i, the probability that a bit sent by the source is received faithfully at the destination is equal to $1 - \varepsilon_i$; this bit is consequently erased with probability ε_i. The elements of the state space $\mathcal{C}$ are indexed in such a way that $i < j$ implies $\varepsilon_i \geq \varepsilon_j$. In its simplest non-trivial instantiation, this finite-state channel with memory can take on two possible states. This specific model is known as the Gilbert-Elliott channel [19], [20], and is illustrated in Fig. 1. Finite-state channels have been employed to model wireless connections in the past, and techniques aimed at selecting representative model parameters are available [24], [25].

The adverse effects of channel uncertainty are countered by the communication system using error-control coding. Data segments are first encoded at the source and then transmitted in the form of codewords to the destination. As discussed above, individual symbols may or may not be received at the destination depending on the realization of the finite-state channel. Decoding is executed on a per codeword basis. We assume that code performance is governed solely by the number of erasures that occur during a transmission interval. Obtaining a distribution for the number of such erasures, conditioned on the initial state of the channel, is therefore highly desirable. There exist various strategies to compute such distributions. Finding the distribution of the channel states and then computing the conditional distribution of the number of erasures is a possible approach [26]. Alternatively, one can employ generating functions and product of matrices with polynomial entries to derive these quantities [27]. At this point, it suffices to point out that

$$\Pr(E = e, C_{N+1} = j | C_1 = i) \tag{2}$$

can be computed efficiently. Above, C_n represents the state of

the channel at time n and E denotes the number of erasures within a block.

We consider a standard and powerful approach to model forward error correction [28]. Every codebook is created using a random binary parity-check matrix $\mathbf{H}$ of size $(N-K) \times N$. The admissible codewords are the elements of the nullspace of $\mathbf{H}$. Decoding at the receiver is executed using a maximum likelihood decision rule. The ensuing probability of decoding failure, conditioned on e erasures, is then given by

$$P_{\mathrm{f}}(N-K, e) = 1 - \prod_{i=0}^{e-1} \left(1 - 2^{i-(N-K)}\right) \quad (3)$$

where N is the code length and K designates the number of information bits per codeword [29]. Accounting for channel states, the conditional probability of decoding failure at the destination, which we represent by $P_{\mathrm{df}}(j; i)$, is given by

$$\sum_{e=0}^{N} P_{\mathrm{f}}(N-K, e) \Pr\left(E = e, C_{N+1} = j | C_1 = i\right).$$

Similarly, the conditional probability of decoding success, labeled $P_{\mathrm{ds}}(j; i)$, can be written as

$$\sum_{e=0}^{N} (1 - P_{\mathrm{f}}(N-K, e)) \Pr\left(E = e, C_{N+1} = j | C_1 = i\right).$$

Thus, combining (2) and (3), we obtain the probability of transitioning to state j with or without decoding success, conditioned on the channel being in state i. Collectively, these probabilities underlie the evolution of the queueing system; this is described next.

To conform with our block encoding scheme, a data packet of length L must be divided into $M = \lceil L/K \rceil$ segments, each of size K. The ending segment of a packet is zero-padded, if needed. Note that M is too a geometric random variable, albeit with parameter

$$\rho_r = \sum_{\ell=1}^{K} (1-\rho)^{\ell-1} \rho = 1 - (1-\rho)^K. \quad (4)$$

These segments are successively encoded into codewords of length N and sent over the finite-state channel. Upon successful decoding, the destination acknowledges reception of the information and the corresponding segment is discarded from the source buffer. On the other hand, when transmission fails, the source is notified. The leading data segment is then immediately re-encoded and transmitted once again over the wireless channel. This process continues until successful reception of the codeword at the destination.

The number of packets awaiting transmission is selected as the state of the queue. This perspective reflects our inclination towards delay-sensitive communications. An alternate formulation would take the number of segments awaiting transmission as the state of the queue. This latter option would be more appropriate to evaluate the size of the memory necessary to store information at the source, at the expense of a less accurate delay characterization. It is worth reemphasizing that, in our framework, a packet departs from the transmit buffer whenever a codeword is decoded successfully at the destination and the corresponding segment is the last parcel of information of the lead packet.

Throughout, we use Q_s to identify the state of the queue at discrete-time s. Although the stochastic process $\{Q_s\}$ does not possess the Markov property, it can be shown that the channel state and the queue length at the onset of a codeword cycle, jointly designated by $Y_s = (C_{sN+1}, Q_s)$, form a Markov chain [17]. The transition probabilities for this Markov chain can be calculated as follows. Suppose that the queue is non-empty, i.e., $Y_s = (i, q)$ where $q > 0$. Owing to our problem definition, the admissible values for Q_{s+1} are $\{q-1, q, q+1\}$. Several factors can affect the evolution of the queue over time: the arrival of a new packet, the successful decoding of a codeword and whether or not this latter codeword is the last segment of a data packet. The only scenario that leads to a decrease in the queue is having no arrival and one packet departure. Recall that a packet departure occurs when a codeword is successfully decoded and the corresponding segment is the last parcel of information of the lead data packet. This yields

$$\mu_{ij} = \Pr(Y_{s+1} = (j, q-1) | Y_s = (i, q))$$
$$= (1-\gamma) P_{\mathrm{ds}}(j; i) \rho_r.$$

For the queue length to remain at a specific level, departures and arrivals must be balanced. In particular, there can be either no departure and no arrival, or one departure and one arrival,

$$\kappa_{ij} = \Pr(Y_{s+1} = (j, q) | Y_s = (i, q))$$
$$= (1-\gamma) \left(P_{\mathrm{df}}(j; i) + P_{\mathrm{ds}}(j; i)(1 - \rho_r)\right)$$
$$+ \gamma P_{\mathrm{ds}}(j; i) \rho_r.$$

Finally, the queue increases whenever a packet arrives and no departure occurs,

$$\lambda_{ij} = \Pr(Y_{s+1} = (j, q+1) | Y_s = (i, q))$$
$$= \gamma \left(P_{\mathrm{df}}(j; i) + P_{\mathrm{ds}}(j; i)(1 - \rho_r)\right).$$

When the queue is empty, $Q_s = 0$, similar arguments apply, except that there can be no departures,

$$\kappa_{ij}^0 = \Pr(Y_{s+1} = (j, 0) | Y_s = (i, 0))$$
$$= (1-\gamma) \Pr(C_{(s+1)N+1} = j | C_{sN+1} = i)$$
$$\lambda_{ij}^0 = \Pr(Y_{s+1} = (j, 1) | Y_s = (i, 0))$$
$$= \gamma \Pr(C_{(s+1)N+1} = j | C_{sN+1} = i).$$

Possible transitions for a non-empty queue at level q are depicted in Fig 2. Again, for simplicity, the diagram assumes a two-state channel at the physical layer.

B. Adaptive Antenna Implementation

Although series of carefully designed experiments in anechoic chambers have been reported previously in the literature on reconfigurable antennas [6], [2], [3], establishing accurate mathematical models for particular system implementations can be a daunting task. Given that this article only offers a preliminary investigation on the topic, we make simplifying assumptions that are somewhat favorable to RF-agile devices. The rational behind this reasoning is to gain insight without
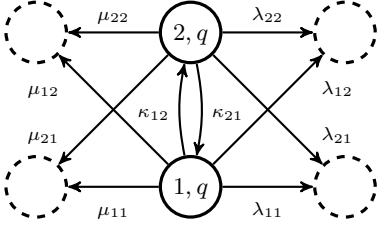
Fig. 2. Possible transitions with partial labeling for a queueing system built upon a two-state channel model and operating using a fixed antenna configuration. Self-transitions are intentionally omitted.

Fig. 3. In this diagram, the two-state antenna system evolves unaltered while in state (c_2, q); whereas an antenna reconfiguration is initiated whenever the system enters state (c_1, q). The reconfiguration process alters the transition probability of the system, as designated by the tildes.

risking to prematurely discard a technology that may eventually lead to significant gains. We postulate that reconfigurable antennas have a large number of possible configurations, and we assume that the wireless channels induced by these configurations are independent from one another. This is more likely to apply to situations where devices are embedded in rich scattering environments. A direct implication of these two hypotheses is that a wireless device equipped with a reconfigurable antenna can always elect to switch to a different virtual channel. Furthermore, once this transformation is accomplished, the probability that the wireless channel occupies a particular state becomes equal to the stationary probability of this same state.

The second aspect of reconfigurable antennas that warrants attention is the latency of the morphing process. The amount of time necessary to execute an antenna reconfiguration depends heavily on the physics underlying the process. Envisioned technologies related to our present investigation include electronic switches, microelectromechanical systems (MEMS) and microfluidic devices. Collectively, these various techniques embody a range of options in terms of latency, efficiency and power consumption. They also offer fundamentally different mechanisms that can provide measurable tradeoffs between speed, power handling, linearity and overall complexity. Moreover, they each feature a compact form factor suitable for mobile devices. Based on the state-of-the-art for these mechanisms, it is reasonable to assume that reconfiguration latency is no greater than a typical codeword cycle (on the order of 4.615 ms). In our analysis, we assume that triggering an antenna reconfiguration event results in the loss of one codeword transmission opportunity; no segment can be decoded at the completion of the corresponding transmission interval and, as such, there cannot be a departure from the queue. This is the price to pay for the opportunity to access a fresh channel realization.

We consider static control policies for antenna handling that are based solely on channel state. Furthermore, we look at hierarchical structures: if i is deemed deficient enough to initiate a channel reconfiguration, then j will also trigger a reconfiguration whenever $j \leq i$. While it may be possible to create idiosyncratic channel probability transition matrices for which this class of policies is highly suboptimal, hierarchical structures are expected to work well for realistic channel models derived from empirical observations. Implicit to these control policies is channel state knowledge at the source; this
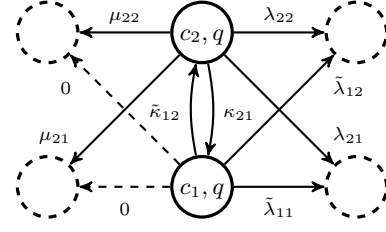
construction again favors adaptive systems. More pragmatic schemes would have to employ state estimates or trigger a reconfiguration based on the number of successive failed decoding attempts. Nevertheless, our simplified framework is a logical starting point: if an adaptive system fails to produce significant gains when side information is available, then there is no need to study more intricate formulations.

When the channel state is deemed satisfactory, no reconfiguration takes place and the transition probabilities defined in Section II-A apply. On the other hand, when a system reconfiguration is initiated, the transition probabilities become

$$\tilde{\mu}_{ij} = \Pr(Y_{s+1} = (j, q-1)|Y_s = (i,q)) = 0$$
$$\tilde{\kappa}_{ij} = \Pr(Y_{s+1} = (j, q)|Y_s = (i,q)) = (1-\gamma)p_C(j)$$
$$\tilde{\lambda}_{ij} = \Pr(Y_{s+1} = (j, q+1)|Y_s = (i,q)) = \gamma p_C(j),$$

where $p_C(\cdot)$ is the marginal probability distribution of the individual wireless channels. We emphasize that these probabilities are completely determined by the latter stationary distribution and the probability of a packet arrival. In a similar fashion, we have

$$\tilde{\kappa}_{ij}^0 = (1-\gamma)p_C(j) \qquad \tilde{\lambda}_{ij}^0 = \gamma p_C(j)$$

whenever a reconfiguration event is sparked from an empty queue. Fig. 3 presents a level-transition diagram for a two-state channel where an antenna reconfiguration is prompted whenever the channel state lies in the lowest level at the onset of a codeword cycle. It may be instructive to compare this graph with Fig. 2, whose labels embody the operation of a communication system with a static antenna structure.

This completes the description of the queueing systems, with and without reconfigurable antenna structures. In both cases, the state space for the discrete-time packetized system is $\mathcal{C} \times \mathbb{N}_0$. Each implementation will be stable provided that the average arrival rate is less than its expected service rate. When this is the case, the underlying Markov chain is positive recurrent and it admits a stationary distribution [30]. In the next section, we further discuss stability conditions and we provide means to compute invariant distributions. This is performed by linking the mathematical formulation of our problem to classical queueing results.

III. QUEUEING BEHAVIOR

Recall that a new data packet arrives at the source at time s with probability γ. Moreover, the number of segments

contained in any packet is a geometric random variable with parameter ρ_r, as described in (4). The expected arrival rate in segments per block is then given by

$$\gamma \mathrm{E}\,[M] = \frac{\gamma}{\rho_r}.$$

The expected service rate depends on the communication scheme employed. In the traditional implementation with static antennas, the progression of the wireless channel is unaltered at the codeword boundaries. The throughput in segments per block can be expressed as

$$\sum_{i \in \mathcal{C}} \sum_{j \in \mathcal{C}} P_{\mathrm{ds}}(j;i) p_C(i) \qquad (5)$$

where $p_C(\cdot)$ represents the stationary channel distribution associated with matrix

$$\mathbf{B}^N = \begin{bmatrix} b_{11}^{(N)} & b_{12}^{(N)} & \cdots & b_{1k}^{(N)} \\ b_{21}^{(N)} & b_{22}^{(N)} & \cdots & b_{2k}^{(N)} \\ \vdots & \vdots & \ddots & \vdots \\ b_{k1}^{(N)} & b_{k2}^{(N)} & \cdots & b_{kk}^{(N)} \end{bmatrix}. \qquad (6)$$

We stress that, since $\mathbf{B}$ is assumed irreducible and aperiodic, its invariant distribution $p_C(\cdot)$ exists and is unique [31]. This distribution is also invariant for probability transition matrix $\mathbf{B}^N$, which justifies its use in (5).

Obtaining the probability transition matrix for the adaptive architecture with reconfigurable antennas is a slightly more involved task. Let $\mathcal{C}^\dagger = \{\ell, \dots, k\}$ represent the collection of channel states judged suitable for data transmission. Then, necessarily, the set $\mathcal{C} \setminus \mathcal{C}^\dagger$ contains all the channel states for which a reconfiguration command is issued. In view of this partitioning, we gather that the probability transition matrix for the channel state at the onset of a codeword cycle is

$$\tilde{\mathbf{B}}^{(N)} = \begin{bmatrix} p_C(1) & p_C(2) & \cdots & p_C(k) \\ \vdots & \vdots & \ddots & \vdots \\ p_C(1) & p_C(2) & \cdots & p_C(k) \\ b_{\ell 1}^{(N)} & b_{\ell 2}^{(N)} & \cdots & b_{\ell k}^{(N)} \\ \vdots & \vdots & \ddots & \vdots \\ b_{k1}^{(N)} & b_{k2}^{(N)} & \cdots & b_{kk}^{(N)} \end{bmatrix}. \qquad (7)$$

We point out that matrix entry $b_{ij}^{(N)}$ is implicitly defined in (6). Given that successful decoding is only possible when a codeword is sent, we can write the throughput for the adaptive system as

$$\sum_{i \in \mathcal{C}^\dagger} \sum_{j \in \mathcal{C}} P_{\mathrm{ds}}(j;i) \tilde{p}_C(i),$$

where $\tilde{p}_C(\cdot)$ is the invariant distribution associated with probability transition matrix $\tilde{\mathbf{B}}^{(N)}$.

When the average arrival rate is strictly less than the expected service rate, Foster's criteria guarantees that the corresponding Markov chain is positive recurrent [30, p. 167]. It is instructive to point out that channel memory and channel quality can greatly influence the expected service rate of a communication system. This phenomenon is rooted in the subtle interactions between channel output sequences and

the probability of decoding success at the destination. It is illustrated through numerical examples in the next section.

The channel state and the queue length at the onset of a codeword cycle, $Y_s = (C_{sN+1}, Q_s)$, jointly form a stochastic process with a countably infinite state space. A natural ordering for its elements is the following,

$$(1,0), \dots, (k,0), (1,1), \dots, (k,1), (1,2), \dots$$

Collectively, the subset of states

$$\{(1,q), \dots, (k,q)\}$$

is known as the qth level of the chain. Using this ordering and the level abstraction, we can introduce a probability transition operator $\mathbf{T}$ for aggregate chain $\{Y_s\}$,

$$\mathbf{T} = \begin{pmatrix} \mathbf{C}_1 & \mathbf{C}_2 & \mathbf{0} & \mathbf{0} & \cdots \\ \mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2 & \mathbf{0} & \cdots \\ \mathbf{0} & \mathbf{A}_0 & \mathbf{A}_1 & \mathbf{A}_2 & \cdots \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_0 & \mathbf{A}_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix}. \qquad (8)$$

In the case of a fixed antenna configuration, the submatrices $\mathbf{A}_0, \mathbf{A}_1, \mathbf{A}_2$ are given by

$$\mathbf{A}_0 = \begin{bmatrix} \mu_{11} & \cdots & \mu_{1k} \\ \vdots & \ddots & \vdots \\ \mu_{k1} & \cdots & \mu_{kk} \end{bmatrix} \qquad \mathbf{A}_1 = \begin{bmatrix} \kappa_{11} & \cdots & \kappa_{1k} \\ \vdots & \ddots & \vdots \\ \kappa_{k1} & \cdots & \kappa_{kk} \end{bmatrix}$$

$$\mathbf{A}_2 = \begin{bmatrix} \lambda_{11} & \cdots & \lambda_{1k} \\ \vdots & \ddots & \vdots \\ \lambda_{k1} & \cdots & \lambda_{kk} \end{bmatrix}.$$

When the queue is empty, the block transitions are governed by

$$\mathbf{C}_1 = \begin{bmatrix} \kappa_{11}^0 & \cdots & \kappa_{1k}^0 \\ \vdots & \ddots & \vdots \\ \kappa_{k1}^0 & \cdots & \kappa_{kk}^0 \end{bmatrix} \qquad \mathbf{C}_2 = \begin{bmatrix} \lambda_{11}^0 & \cdots & \lambda_{1k}^0 \\ \vdots & \ddots & \vdots \\ \lambda_{k1}^0 & \cdots & \lambda_{kk}^0 \end{bmatrix}.$$

These definitions have analog counterparts for systems with reconfigurable antennas; the appropriate entries are simply replaced by their homologues,

$$\tilde{\mu}_{ij}, \tilde{\kappa}_{ij}, \tilde{\lambda}_{ij}, \tilde{\kappa}_{ij}^0 \text{ and } \tilde{\lambda}_{ij}^0.$$

In both scenarios, with and without adaptation, the ensuing Markov chains belong to the class of random processes with repetitive structures [22]. One can find the stationary distribution associated with $\mathbf{T}$ by inspecting the substochastic matrix $\mathbf{U}$ whose entry u_{ij} denotes the probability that, starting from state $(1,i)$, the Markov chain Y_s first re-enters level one through $(1,j)$ and does so without visiting any state at level zero. A probabilistic path-counting argument leads to Proposition 3.1.

Proposition 3.1: Define $\mathbf{U}_1 = \mathbf{A}_1$. The iterative expression

$$\mathbf{U}_{m+1} = \mathbf{A}_1 + \mathbf{A}_2 \left(\mathbf{I} - \mathbf{U}_m \right)^{-1} \mathbf{A}_0$$

is well-defined for all $m \in \mathbb{N}$; its limit exists and

$$\lim_{m \to \infty} \mathbf{U}_m = \mathbf{U}.$$

Let π represent the invariant distribution of the augmented Markov chain, and denote the subcomponents associated with level q by π_q,

$$\pi_q = (\Pr(Y = (1, q)), \ldots, \Pr(Y = (k, q))).$$

Proposition 3.2: Define $\mathbf{R} = \mathbf{A}_2(\mathbf{I} - \mathbf{U})^{-1}$ and recall that the entries of π are non-negative and sum up to one. The invariant distribution induced by $\mathbf{T}$ is entirely determined through the following relations,

$$\begin{bmatrix} \pi_0 & \pi_1 \end{bmatrix} \begin{bmatrix} \mathbf{C}_1 & \mathbf{C}_2 \\ \mathbf{A}_0 & \mathbf{A}_1 + \mathbf{R}\mathbf{A}_0 \end{bmatrix} = \begin{bmatrix} \pi_0 & \pi_1 \end{bmatrix} \qquad (9)$$

and $\pi_q = \pi_1 \mathbf{R}^{q-1}$ for $q \geq 1$.

Propositions 3.1 & 3.2 provide an algorithmic blueprint to compute the stationary distribution of the augmented Markov chain: obtain $\mathbf{U}$ through repeated iterations; compute $\mathbf{R}$ and form irreducible and aperiodic probability transition matrix

$$\begin{bmatrix} \mathbf{C}_1 & \mathbf{C}_2 \\ \mathbf{A}_0 & \mathbf{A}_1 + \mathbf{R}\mathbf{A}_0 \end{bmatrix};$$

find its invariant distribution; append missing values of π using $\pi_q = \pi_1 \mathbf{R}^{q-1}$ and normalize.

IV. PERFORMANCE EVALUATION

Once the stationary distribution is acquired, we can compute several performance criteria of interest. In this article, we examine the average delay, the probability that the queue length exceeds a certain threshold and the decay rate of the queue occupancy. First, we note that the expected queue length is given by

$$\sum_{q=0}^{\infty} q\pi_q \cdot \mathbf{1} = \pi_1 \left(\sum_{q=1}^{\infty} q\mathbf{R}^{q-1} \right) \cdot \mathbf{1}, \qquad (10)$$

where $\mathbf{1}$ is a column vector of all ones. Using Little's formula, we deduce that the mean waiting time in the queue is simply (10) divided by expected arrival rate γ. The decay rate of the queue occupancy can be written as

$$\lim_{\tau \to \infty} \frac{1}{\tau} \log \Pr(Q \geq \tau) = \log \varrho(\mathbf{R}), \qquad (11)$$

where $\varrho(\mathbf{R})$ is the spectral radius of $\mathbf{R}$; and its complementary cumulative distribution function is determined by the finite sum

$$\Pr(Q > \tau) = 1 - \sum_{q=0}^{\lfloor \tau \rfloor} \pi_q \cdot \mathbf{1}.$$

In the next section, we use these performance criteria to show that time-dependencies in the underlying physical channel can adversely affect the behavior of a queueing system. Moreover, having the ability to reconfigure an antenna structure at appropriate moments can help mitigate the undesirable effects of channel memory.

A. *Partially Observable Systems*

Up to this point, we have looked at the possible benefits of reconfigurable antenna systems by comparing the performance of specific policies using various optimization objectives, e.g., throughput, average delay, queue threshold violation. To offer a different perspective, we study the communication process in a decision theoretic framework. At the onset of every codeword, the transmitter has the freedom to select an appropriate code-rate or it can trigger an antenna reconfiguration when the channel quality is deemed deficient. The exact state of the channel is hidden from the transmitter, yet the device can estimate it from the acknowledgements of previous transmission attempts. In this latter scenario, the objective is to maximize the amount of information transmitted in the long run. The formulation of a *partially observable Markov decision process* (POMDP) is immediate. A distinguishing feature of the proposed framework is that it naturally incorporates state estimation, code rate selection, antenna reconfiguration and throughput maximization.

We note that the transmitter must base its actions on past observations. To act optimally, it is sufficient, surprisingly, to summarize previous observations into a *belief state* [32]. One such choice for the belief state is the probability distribution of the channel states conditioned on the past observations. Whenever the transmitter takes an action and collects an observation, this belief state must be updated. We refer the reader to [33] for the notation and solution methodologies for POMDPs. Below, we formally define the quantities involved in our problem formulation.

A partially-observable Markov decision process consists of a tuple $(\mathcal{C}, \mathcal{B}, \mathcal{A}, \mathcal{O}, \mathcal{T}, \mathcal{R}, \Omega)$, as described below.

1) As before, $\mathcal{C}$ represents the set of admissible channel states, $\{1, \ldots, k\}$, and C_n denotes the realization of the channel at time instant n.
2) $\mathcal{B} \subset [0,1]^k$ is the collection of prior distributions (or *belief states*) on the channel states at the transmitter. The belief at time instant n is written as B_n.
3) $\mathcal{A}$ is the space of actions, which includes different code rates and the reconfiguration event. Code rates of the form a/N, $1 \leq a \leq N$, $a \in \mathbb{N}$ are valid options. We index the code-rate a/N with a, and a reconfiguration event is labeled by 0. The action at time instant n is denoted by A_n.
4) $\mathcal{O}$ denotes the set of observations; these assume the form of an ACK or NACK obtained form the receiver. The observation at time instant n is denoted by O_n.
5) $\Omega : \mathcal{B} \times \mathcal{A} \times \mathcal{O} \to \mathbb{R}^+$ is the observation function which gives, for each action and a belief state, the probability of making an observation.
6) $\mathcal{T} : \mathcal{B} \times \mathcal{A} \times \mathcal{O} \to \mathcal{B}$ denotes the transition function which produces, for each action and an observation, the updated belief state as a function of the previous belief state.
7) $\mathcal{R} : \mathcal{B} \times \mathcal{A} \to \mathbb{R}^+$ gives the reward received for the corresponding action. For our problem, this is the number of transmitted information bits, if any, weighted by the probability of the successful reception.

Having established a proper framework, we proceed with the computation of the quantities $\mathcal{T}$, $\mathcal{R}$, Ω for our problem. By definition, the observation function is given by

$$\Omega(\psi, a, o) = \Pr(O_{n+1} = o | B_n = \psi, A_n = a),$$

and, hence,

$$\Omega(\psi, a, o) = \begin{cases} \psi \mathbf{A}_0 \mathbf{1}, & o = \text{ACK} \\ \psi \mathbf{A}_1 \mathbf{1}, & o = \text{NACK} \end{cases}$$

for $a \neq 0$. When a reconfiguration request is issued, we invariably get $\Omega(\psi, 0, \text{ACK}) = 0$ and $\Omega(\psi, 0, \text{NACK}) = 1$. The transfer function $\mathcal{T}$ updates the current belief state based on the action taken and the observation made. That is,

$$[\mathcal{T}(\psi, a, o)]_i = \Pr(C_{n+1} = i | B_n = \psi, O_n = o, A_n = a).$$

This can be reduced to

$$[\mathcal{T}(\psi, a, o)]_i = \frac{\Pr(C_{n+1} = i, O_n = o | A_n = a, B_n = \psi)}{\Pr(O_n = o | A_n = a, B_n = \psi)}.$$

Writing the above equation compactly, for $a \neq 0$,

$$\mathcal{T}(\psi, a, \text{ACK}) = \frac{\psi \mathbf{A}_0}{\psi \mathbf{A}_0 \mathbf{1}}, \quad \mathcal{T}(\psi, a, \text{NACK}) = \frac{\psi \mathbf{A}_1}{\psi \mathbf{A}_1 \mathbf{1}}.$$

For the reconfiguration event, which we assume leads to a new virtual channel, we obtain

$$\mathcal{T}(\psi, 0, \text{ACK}) = \mathcal{T}(\psi, 0, \text{NACK}) = p_C.$$

The reward function is the expected number of information bits received at the destination under state ψ upon action a, normalized by the length of a codeword,

$$\mathcal{R}(\psi, a) = \tfrac{a}{N} \psi \mathbf{A}_0 \mathbf{1}.$$

We define an n-step policy $\Delta_n = \{\delta_1, \ldots, \delta_n\}$ as a collection of n functions that map each belief state in $\mathcal{B}$ to an action from $\mathcal{A}$. The expected discounted reward function $V_{\Delta_t, t} : \mathcal{B} \to \mathbb{R}^+$ is given by

$$V_{\Delta_n, n}(\psi) = \mathrm{E}\left[\sum_{t=1}^n \beta^{t-1} \mathcal{R}(B_t, \delta_t(B_t))\right],$$

where $B_1 = \psi$. Intuitively, $V_{\Delta_n, n}$ denotes the expected rewards received by executing the policy Δ_n for n-steps. However, the rewards procured in the future are discounted by a factor of β each time. In a delay-sensitive system, the discount factor can be thought of as penalizing the information bits that will be transmitted in the future. Still, this is a standard approach that ensures convergence and the existence of a solution. The n-step finite horizon value function is defined as the maximum of $V_{\Delta_n, n}$ over all the policies,

$$V_n = \max_{\Delta_n} V_{\Delta_n, n}.$$

By the *principle of optimality*, the finite horizon value functions satisfy [34]

$$V_n(\psi) = \max_{a \in \mathcal{A}} \Big[\mathcal{R}(\psi, a) + \beta \sum_{o \in \mathcal{O}} \Omega(o, \psi, a) V_{n-1}(\mathcal{T}(\psi, o, a)) \Big].$$

Similarly, the infinite horizon value function is defined as

$$V(\psi) = \max_{\Delta} \mathrm{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \mathcal{R}(B_t, \delta_t(B_t))\right],$$

where $B_1 = \psi$ and Δ is a sequence of functions $\{\delta_t\}_{t=1}^{\infty}$ which map each belief state to an action. The optimal policy Δ^* that maximizes the above equation is a stationary policy [35]. That is $\Delta^* = \{\delta_t = \delta^*\}_{t=1}^{\infty}$ for some function $\delta^* : \mathcal{B} \to \mathcal{A}$. The infinite horizon value function can be approximated by a finite horizon value function according to [35, Theorem 3]

$$V(\psi) = \lim_{n \to \infty} V_n(\psi).$$

V. NUMERICAL RESULTS

Although the methodology introduced above applies to general finite-state erasure channels with memory, our numerical study focuses on the Gilbert-Elliott model depicted in Fig. 1. This model has received a fair amount of attention in the literature. It captures many of the features associated with wireless environments such as uncertainty, fading and channel memory. Yet, it remains simple enough for a straightforward exposition of our findings. Overall, the Gilbert-Elliott model provides valuable insights about the operation of wireless communication systems without being overly intricate, which gives ground for its adoption. This is especially relevant for a first characterization of the potential benefits associated with reconfigurable antenna structures.

The choice of a hierarchical control policy for the Gilbert-Elliott channel with side information is straightforward. The only non-trivial candidate is the adaptive policy where the source triggers an antenna reconfiguration whenever the state of the channel is 1 at the onset of a codeword transmission. The performance of this adaptive scheme is compared with the operation of a static system where the antenna structure is fixed and codewords are sent at every opportunity.

For the Gilbert-Elliott channel model, the stochastic matrix $\mathbf{B}$ reduces to

$$\mathbf{B} = \begin{bmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{bmatrix},$$

and it has only two degrees of freedom. One possible way to portray these degrees of freedom is to talk about the stationary probabilities of the states,

$$\Pr(C = 1) = \frac{b_{21}}{b_{12} + b_{21}} \qquad \Pr(C = 2) = \frac{b_{12}}{b_{12} + b_{21}},$$

and channel memory; this is the approach we use throughout. The memory of a Gilbert-Elliott channel can be expressed at the symbol level using $1 - b_{12} - b_{21}$. An equivalent way to characterize memory is to consider changes at the codeword level,

$$1 - b_{12}^{(N)} - b_{21}^{(N)} = (1 - b_{12} - b_{21})^N \in [0, 1). \tag{12}$$

It is typically more insightful to plot results using the latter scaling and, as such, this is the unit we employ in our figures. Additional system parameters are selected to approximate the operation of a GSM communication link. The block length is set to $N = 114$. New packets arrive at the source with

probability $\gamma = 0.20$, and their expected length is $\rho^{-1} = 195$ bits. Based on a 4.615 ms codeword cycle, this yields a lightly loaded connection at roughly 8.45 kbps; these are realistic numbers for digital telephony.

We assess the performance of our competing systems when operating over erasure channels with an erasure probability equal to 20 percent. We first explore the impact of channel memory on overall performance. We study a channel model with $b_{12} = 4b_{21}$, $\varepsilon_1 = 0.5$ and $\varepsilon_2 = 0.125$. Channel correlation over time is varied progressively from the memoryless case to a very slow fading profile. Note that the value of K is throughput optimized for every parameter set and system implementation, leading to a fair comparison between schemes.

Figure 4 displays maximum throughput in bits per channel use as a function of the memory coefficient defined in (12). When channel memory is small, the communication system
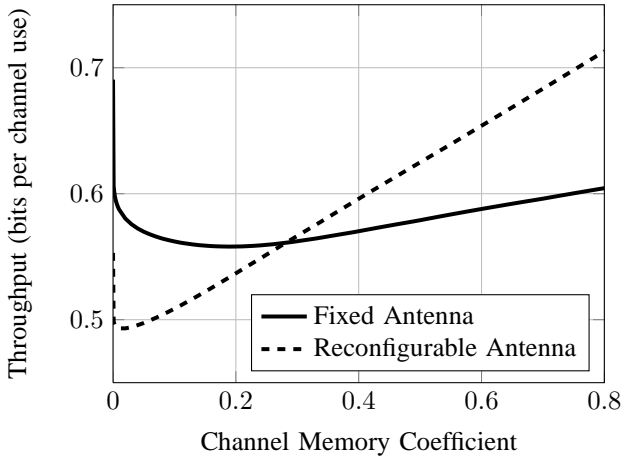

Fig. 4. The throughput of a traditional communication system is compared to that of a competing implementation with a reconfigurable antenna. As channel memory increases, the performance of the reconfigurable system surpasses the maximum service rate of the traditional implementation.

with a fixed antenna structure performs better. In particular, if the mixing time of the Gilbert-Elliott channel is shorter than a codeword transmission cycle, then reconfiguration offers little rewards. It is therefore more profitable to send codewords constantly. On the other hand, as the memory coefficient approaches one, the channel can get stuck in a bad fade for a prolonged period of time. This phenomenon almost certainly guarantees decoding failure at the next attempt. This prompts the RF front-end to trigger a reconfiguration and seek a more auspicious channel realization. The crossover point in Fig. 4 is approximately 0.28. Interestingly, at this crossover point, the expected sojourn time in state 1 is approximately 113 bits, which is very close to the actual block length.

Similar curves can be generated for other parameter sets. Each of these curves identifies a crossover point in terms of channel memory and erasure probability where the throughput of the adaptive system overtakes the expected service rate of the traditional implementation. Plotting these points delineates the boundary of two regions, one where a static system with a fixed antenna structure performs better and

a second region where the reconfigurable implementation delivers enhanced performance. This is illustrated in Fig. 5. We
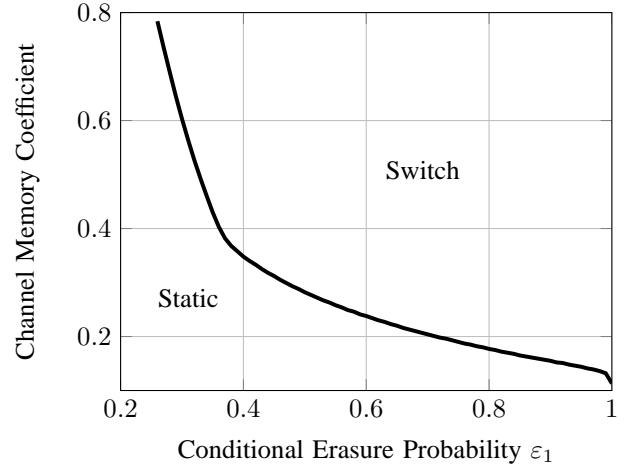

Fig. 5. System parameters determine which implementation delivers better throughput, the static system with a fixed antenna structure (Static) or the switching scheme with reconfigurable antennas (Switch). Correlation and fade differentiation are advantageous to the RF-agile switching scheme. In this figure, the probability of erasure is set at 0.20 and $\varepsilon_2 = (1 - \varepsilon_1)/4$.

immediately see from this figure that channel correlation over time favors reconfigurable antennas. In addition, experiencing vastly different channel qualities over the various fade levels benefits adaptive implementations. Altogether, the capabilities of reconfigurable antenna systems seem better suited to harsh wireless environments.

We supplement the preceding results by investigating the performance of the two competing systems in relation to mean waiting time. This latter criterion is appropriate for lightly loaded connections and delay-sensitive applications such as mobile telephony and video conferencing. Again, we maintain the average rate of bit erasure at 20 percent, and we set the conditional probabilities of erasure to $\varepsilon_1 = 0.5$ and $\varepsilon_2 = 0.125$. As before, we vary the channel memory coefficient to first produce a memoryless process followed by increasingly correlated erasure sequences. Figure 6 plots the mean waiting time at the source as a function of memory. The crossover point where the switching system with a reconfigurable antenna structure overtakes the static implementation is approximately the same as in the case of throughput. In fact, preliminary results indicate that similar behavior can be observed for various parameter sets and different optimization criteria including mean waiting time, asymptotic decay rate in queue occupancy and threshold violation probability. This robustness may be attributable to the simplicity of the Gilbert-Elliott model and may not hold for more complex channel models. This observation warrants further research, but this lies outside the scope of our article. In practice, this suggests that good performance can be achieved with RF-agile antenna structures by identifying regions where reconfiguration should take place. The system can then estimate the current state of the channel and decide, according to its local map, whether or not a reconfiguration event should be triggered.

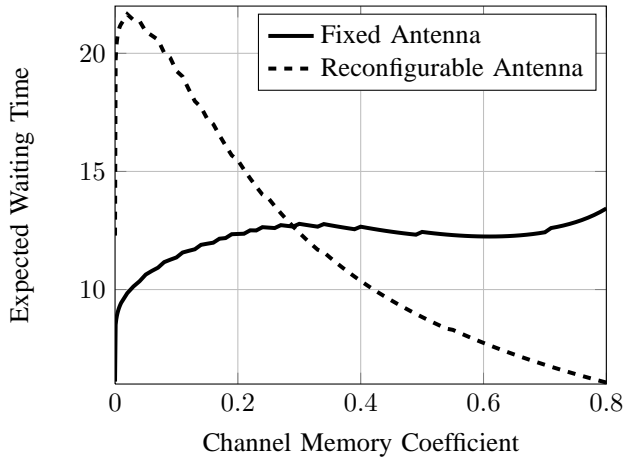Fig. 7 shows the numerical analysis of a two-state channel

Fig. 6. Mean waiting times for traditional and switching systems are plotted as functions of channel memory; a smaller waiting time is desirable. When the channel is weakly correlated over time, the system with a fixed antenna configuration performs better. On the other hand, in slow fading scenarios, the adaptive implementation with a reconfigurable antenna structure becomes advantageous.
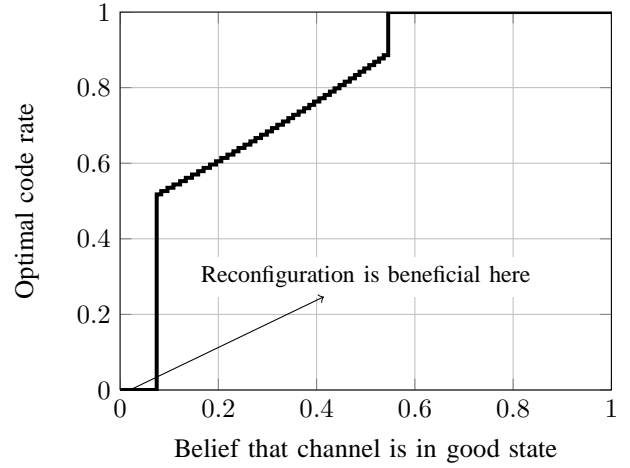

Fig. 7. This figure shows the choice of optimal code rate versus the belief that the channel is in good state. Consistent with intuition, when the transmitter is more confident that the channel is in good state, it can use a high rate code. Conversely, when the transmitter is certain that the channel is in bad state, reconfiguring the antenna is beneficial. We note that the optimal code rate is a threshold based policy, partitioning the belief space.

from the POMDP perspective presented in Section IV-A. The parameters for the analysis are: $\varepsilon_1 = 1$, $\varepsilon_2 = 0$, a channel memory coefficient of 0.3, an average erasure rate of 20%, and discount factor $\beta = 0.9$. For this two-state channel, the belief vector lies in a two-dimensional simplex. Yet, since the two elements in the belief vector sum to one, only the second coefficient is kept, namely the belief that the channel is in a good state. More precisely, Fig. 7 shows the decision regions (optimal action to choose i.e., δ^*) as a function of the belief that the channel is in a good state. We observe that the optimal policy is a threshold based policy partitioning the belief space. The optimal code rate is monotonically increasing as a function of belief state. Fig. 8 shows a similar analysis for a three-state channel with a transition probability matrix

$$\mathbf{B} = \begin{bmatrix} 0.998 & 0.002 & 0 \\ 0.001 & 0.998 & 0.001 \\ 0 & 0.002 & 0.998 \end{bmatrix},$$

and erasure values $\varepsilon_1 = 1$, $\varepsilon_2 = 0.15$, $\varepsilon_3 = 0$. This gives an average erasure rate of 20% and a channel memory coefficient of 0.8. The discount factor used for the computation of infinite horizon value function is $\beta = 0.9$. Fig. 9 shows the expected value function (averaged over the belief space) for a two-state channel with $\varepsilon_1 = 1$, $\varepsilon_2 = 0$, an average erasure rate of 20%, and discount factor $\beta = 0.9$. In the high memory regime, whenever the channel is in bad state, successful receptions in the fixed antenna system are heavily delayed. Since, the delayed rewards are discounted by a factor of $\beta = 0.9$, the average value function is small in this regime for this system. On the contrary, in the reconfigurable antenna system, prolonged bad channel states can be circumvented by issuing a reconfiguration event and thereafter transmitting during the prolonged good channel states, thus exploiting the benefits of high channel memory. This explains the large gap between the mean value functions of the two systems in this regime.
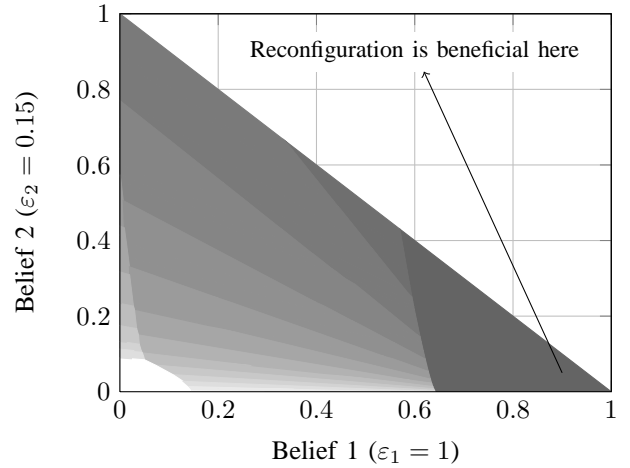

Fig. 8. This figure shows the choice of optimal code rate according to the belief. The darkest region corresponds to the belief space where a reconfiguration event is beneficial. The region in white corresponds to the belief space where using a code rate of 1 is beneficial. The regions in gray correspond to different code rates; the lighter the region, the higher the code rate.

VI. CONCLUDING REMARKS

This preliminary study offers supporting evidence to the claim that reconfigurable antenna structures can improve the performance of communication systems significantly. For the reconfiguration process to be beneficial, the potential rewards of a reconfiguration event must offset the costs of a loss of a codeword transmission opportunity. Two conditions appear to influence this balance. The coherence time of the physical channel must be on the order of the codeword cycle or longer. Furthermore, the quality of the channel must vary significantly over the different fade levels. Slow fading channels appear to be great prospects for reconfigurable antenna systems with adaptive control policies.
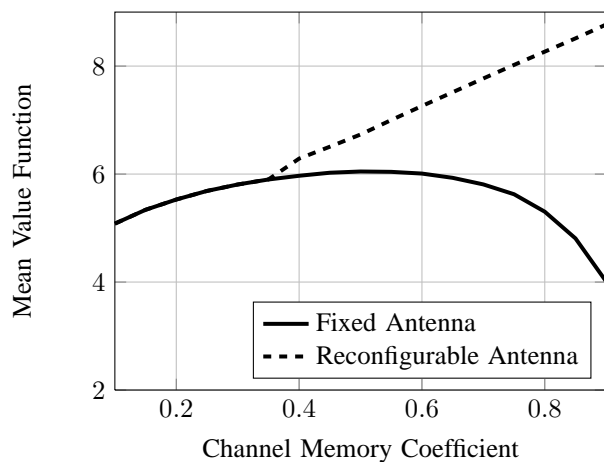
Fig. 9. Comparing the mean value functions for both traditional system and reconfigurable antenna system. The greater flexibility of the adaptive system yields a dominating performance curve. Still, adaptation only pays off in harsh environments where performance is strictly better.

Future studies should address practical issues. Once side information becomes available at the transmitter, power control and scheduling can be employed in conjunction with reconfigurable antenna structures. Extending the queueing formulation to account for these techniques is an interesting goal. Also, the postulate that virtual channels are independent from one another should be explored through empirical measurements. A strong positive correlation among virtual channels would reduce the expected returns of a reconfiguration event. These are promising avenues of future research that may broaden the application potential of reconfigurable antennas and help improve the performance of wireless communication systems.

REFERENCES

[1] G. H. Huff, J. Feng, S. Zhang, G. Cung, and J. T. Bernhard, "Directional reconfigurable antennas on laptop computers: Simulation, measurement and evaluation of candidate integration positions," *IEEE Transactions on Antennas and Propagation*, vol. 52, no. 12, pp. 3220–3227, 2004.

[2] T. L. Roach, G. H. Huff, and J. T. Bernhard, "A comparative study of diversity gain and spatial coverage: Fixed versus reconfigurable antennas for portable devices," *Microwave and Optical Technology Letters*, vol. 49, no. 3, pp. 535–539, 2007.

[3] S. Yang, C. Zhang, H. Pan, A. Fathy, and V. Nair, "Frequency-reconfigurable antennas for multiradio wireless platforms," *IEEE Microwave Magazine*, vol. 10, no. 1, pp. 66–83, 2009.

[4] H. Rajagopalan, Y. Rahmat-Samii, and W. A. Imbriale, "RF MEMS actuated reconfigurable reflectarray patch-slot element," *IEEE Transactions on Antennas and Propagation*, vol. 56, no. 12, pp. 3689–3699, 2008.

[5] G. H. Huff, D. L. Rolando, P. Walters, and J. McDonald, "A frequency reconfigurable dielectric resonator antenna using colloidal dispersions," *IEEE Antennas and Wireless Propagation Letters*, vol. 9, pp. 288–290, 2010.

[6] C. J. Panagamuwa, A. Chauraya, and J. C. Vardaxoglou, "Frequency and beam reconfigurable antenna using photoconducting switches," *IEEE Transactions on Antennas and Propagation*, vol. 54, no. 2, pp. 449–454, 2006.

[7] D. Piazza, N. J. Kirsch, A. Forenza, R. W. Heath, and K. R. Dandekar, "Design and evaluation of a reconfigurable antenna array for MIMO systems," *IEEE Transactions on Antennas and Propagation*, vol. 56, no. 3, pp. 869–881, 2008.

[8] A. Mahanfar, C. Menon, and R. G. Vaughan, "Smart antennas using electro-active polymers for deformable parasitic elements," *Electronics Letters*, vol. 44, no. 19, pp. 1113–1114, 2008.

[9] M. P. Daly and J. T. Bernhard, "Beamsteering in pattern reconfigurable arrays using directional modulation," *IEEE Transactions on Antennas and Propagation*, vol. 58, no. 7, pp. 2259–2265, 2010.

[10] A. F. Molisch and M. Z. Win, "MIMO systems with antenna selection," *IEEE Microwave Magazine*, vol. 5, no. 1, pp. 46–56, 2004.

[11] T. Gou, C. Wang, and S. A. Jafar, "Aiming perfectly in the dark – blind interference alignment through staggered antenna switching," *IEEE Transactions on Signal Processing*, vol. 59, no. 6, pp. 2734–2744, 2011.

[12] J. G. Kim and M. M. Krunz, "Bandwidth allocation in wireless networks with guaranteed packet-loss performance," *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, pp. 337–349, June 2000.

[13] R. Berry and R. Gallager, "Communication over fading channels with delay constraints," *IEEE Transactions on Information Theory*, vol. 48, no. 5, pp. 1135–1149, May 2002.

[14] D. Wu and R. Negi, "Effective capacity: a wireless link model for support of quality of service," *IEEE Transactions on Wireless Communications*, vol. 2, no. 4, pp. 630–643, 2003.

[15] A. Laourine and L. Tong, "Betting on Gilbert-Elliot channels," *IEEE Transactions on Wireless Communications*, vol. 9, no. 2, pp. 723–733, 2010.

[16] L. Liu, P. Parag, J. Tang, W.-Y. Chen, and J.-F. Chamberland, "Resource allocation and quality of service evaluation for wireless communication systems using fluid models," *IEEE Transactions on Information Theory*, vol. 53, no. 5, pp. 1767–1777, May 2007.

[17] P. Parag, J.-F. Chamberland, H. D. Pfister, and K. R. Narayanan, "On the queueing behavior of random codes over a Gilbert-Elliott erasure channel," in *Information Theory Proceedings (ISIT), 2010 IEEE International Symposium on*. IEEE, 2010, pp. 1798–1802.

[18] P. Parag, J. Chamberland, H. Pfister, and K. Narayanan, "Code-rate selection, queueing behavior and the correlated erasure channel," *IEEE Transactions on Information Theory*, 2012, to appear.

[19] E. N. Gilbert, "Capacity of a burst-noise channel," *Bell Syst. Tech. J*, vol. 39, no. 5, pp. 1253–1265, 1960.

[20] E. O. Elliott, "Estimates of error rates for codes on burst-noise channels," *Bell Syst. Tech. J*, vol. 42, no. 5, pp. 1977–1997, 1963.

[21] M. F. Neuts, *Structured stochastic matrices of M/G/1 type and their applications*. CRC, 1989.

[22] ——, *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Dover Publications, 1981.

[23] B. Hajek, "Birth-and-death processes on the integers with phases and general boundaries," *Journal of Applied Probability*, vol. 19, no. 3, pp. 488–499, 1982.

[24] H. S. Wang and N. Moayeri, "Finite state Markov channel – A useful model for radio communication channels," *IEEE Transactions on Vehicular Technology*, vol. 44, no. 1, pp. 163–171, February 1995.

[25] P. Sadeghi, R. A. Kennedy, P. Rapajic, and R. Shams, "Finite-state Markov modeling of fading channels: A survey of principles and applications," *IEEE Signal Processing Magazine*, vol. 25, no. 5, pp. 57–80, September 2008.

[26] L. Wilhelmsson and L. B. Milstein, "On the effect of imperfect interleaving for the Gilbert-Elliott channel," *IEEE Transactions on Communications*, vol. 47, no. 5, pp. 681–688, May 1999.

[27] J.-F. Chamberland, H. D. Pfister, and S. Shakkottai, "First-passage time analysis for digital communication over erasure channels with delay-sensitive traffic," in *Allerton Conference on Communication, Control, and Computing*. Monticello, IL: IEEE, September 2010.

[28] R. G. Gallager, *Information Theory and Reliable Communication*. Wiley, 1968.

[29] T. Richardson and R. Urbanke, *Modern Coding Theory*. Cambridge University Press, 2008.

[30] P. Brémaud, *Markov Chains: Gibbs Fields, Monte Carlo Simulation, and Queues*, ser. Texts in Applied Mathematics. Springer Verlag, 1999, vol. 31.

[31] J. R. Norris, *Markov Chains*, ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.

[32] R. Smallwood and E. Sondik, "The optimal control of partially observable markov processes over a finite horizon," *Operations Research*, pp. 1071–1088, 1973.

[33] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, "Planning and acting in partially observable stochastic domains," *Artificial Intelligence*, no. 1-2, pp. 99–134, May 1998.

[34] D. Bertsekas, *Dynamic Programming and Optimal Control: Volume I*, ser. Athena Scientific optimization and computation series. Athena Scientific, 2005, vol. 1.

[35] E. Sondik, "The optimal control of partially observable markov processes over the infinite horizon: Discounted costs," *Operations Research*, pp. 282–304, 1978.